

## PATENT APPLICATION

This application is submitted in the name of the following inventor:

<i>Inventor</i>	<i>Citizenship</i>	<i>Residence City and State</i>
Keith SON	United States	Hayward, California

The assignee is *Network Appliance*, a corporation having an address in Sunnyvale, California.

## TITLE OF THE INVENTION

Fiber Channel Adaptor for Serial or Parallel ATA Disks

## BACKGROUND OF THE INVENTION

1. *Field of the Invention*

The invention relates to mass storage devices.

2. *Related Art*

It is often useful for mass storage, including disk drives and disk drive subsystems, to be coupled to more than one access point, such as for example more than one computer system, file server, caching device, or other system used for accessing that mass storage. For example, it is often useful for mass storage systems to be redundantly coupled to provide failover or load balancing, or to be multiply connected in the service of providing clustering, high availability, or hot swapping. Known systems for multiply coupling mass storage include SCSI (small computer systems interface) and FC-AL (fiber channel arbitrated loop). While these known systems are relatively robust in their ability to provide multiple coupling for mass storage systems, they have the drawback that they are relatively expensive and have relatively lower storage density per unit cost.

One problem with the known art is that some mass storage systems, including those mass storage devices and subsystems with the least cost and greatest storage density per unit cost, are not well-suited for being coupled to more than one controller. Examples include ATA disks, such as for example S-ATA (serial ATA) disk drives and P-ATA (parallel ATA) disk drives. Accordingly, it would be advantageous to be able to couple such mass storage systems to more than one controlling system. Moreover, it would also be advantageous to be able to use systems or techniques that are available for use with SCSI or FC-AL to provide robust multiple connectivity, with

1 the effect of allowing such mass storage systems (for example, relatively inexpensive  
2 ATA disk drives) with systems and techniques for failover, load balancing, clustering,  
3 high availability, and hot swapping.

4  
5 Accordingly, it would be advantageous to provide techniques by which  
6 disk drives having only a single controller connection might be connected to more than  
7 one controlling device.

#### 8 9 SUMMARY OF THE INVENTION

10  
11 The invention provides a method and system capable of coupling disk  
12 drives having only a single controller connection to more than one controlling device.  
13 In one embodiment, an adaptor provides an ATA disk drive with at least one interface,  
14 and preferably two independent interfaces, to a connector substantially like a fiber  
15 channel backplane (or another type of backplane, not necessarily limited to fiber chan-  
16 nel type backplanes), and having a form factor substantially like a fiber channel back-  
17 plane interface (or another type of backplane, not necessarily limited to fiber channel  
18 type backplanes). The two independent interfaces are capable of being controlled by  
19 switching logic, with the effect that the ATA disk drive is provided with redundant ca-  
20 pability for connecting to controlling devices.

1           In an aspect of the invention, the switching logic is capable of determining  
2   which one of the two independent interfaces is used by the ATA disk drive, with the ef-  
3   fect that the ATA disk drive can be redundantly coupled to a fiber channel loop, or any  
4   other type of redundant storage access structure.

5  
6           In an aspect of the invention, the fiber channel (or other) backplane is ca-  
7   pable of communicating signals suitable to the ATA disk drive, while having an inter-  
8   face to the ATA disk drive substantially with the same form factor as for a fiber channel  
9   disk drive (or other disk drive suited to the backplane). In alternative embodiments, the  
10   fiber channel (or other) backplane is capable of communicating signals suitable to an  
11   industry-standard fiber channel (or other) backplane, with the effect that ATA disk  
12   drives and FC disk drives can be coupled to the same fiber channel backplane.

13  
14           In an aspect of the invention, the fiber channel (or other) backplane is  
15   coupled to an ATA/FC signal converter, with the effect that an ATA disk drive can be  
16   coupled to an interface designed to couple with a fiber channel disk drive (or other disk  
17   drive suited to the backplane). Similarly, the ATA/FC signal converter provides the  
18   capability for fiber channel (or other) backplanes coupled to ATA disk drives to be  
19   mixed with fiber channel (or other) backplanes coupled to fiber channel disk drives (or  
20   other disk drives suited to the backplane), on the identical fiber channel loop (or an-  
21   other redundant storage access structure).

1           One advantage provided by these aspects of the invention is that ATA  
2 disk drives (and other non-FC disk drives) can be coupled to fiber channel (or other)  
3 backplanes, with the effect that industry-standard (and nonstandard) fiber channel (or  
4 other) techniques can be used to provide multiple controllers for such disk drives, and  
5 to provide advantages allowed by such multiple controllers (such as for example,  
6 failover, load balancing, clustering, high availability, and hot swapping).

7  
8           A second advantage provided by these aspects of the invention is that  
9 disk drives of differing types (such as for example P-ATA and S-ATA disk drives) can  
10 be coupled to a common storage access structure, with the effect that P-ATA and S-ATA  
11 disk drives (a) can be used concurrently without inducing substantial additional com-  
12 plexity, and (b) can have data transferred there-between without specific hardware for  
13 converting between parallel-to-serial or serial-to-parallel disk drive interfaces.

14  
15           A third advantage provided by these aspects of the invention is that disk  
16 drives of differing types (such as for example ATA and fiber channel disk drives) can be  
17 coupled to a common storage access structure, even when those disk drives have sub-  
18 stantially different form factors and interfaces, with the effect that a common form fac-  
19 tor and interface (such as for example a fiber channel disk drive form factor and inter-  
20 face) can be used in a storage access structure. Use of a common form factor and inter-  
21 face can reduce physical and other stresses on the actual disk drives included in the  
22 housing.

1  
2           In preferred embodiments, the system includes (1) a housing for an ATA  
3 disk drive, such as for example a P-ATA or an S-ATA disk drive, capable of allowing an  
4 ATA drive to be fitted to another type of form factor or interface, such as for example a  
5 fiber channel drive form factor and interface, (2) an adaptor capable of coupling the  
6 ATA disk drive to a backplane, such as for example a fiber channel backplane, the  
7 adaptor having two independent interfaces, and (3) first switching logic capable of se-  
8 lecting between the two independent interfaces, with the effect that the ATA disk drive  
9 can be redundantly coupled to a storage system. The first switching logic might include  
10 hardware, firmware, software, or a combination or mixture thereof, and might be pro-  
11 grammable. The first switching logic might also include control of a switched signal  
12 capable of supplying power to the disk drive.

13  
14           In such embodiments, the system preferably also includes (4) converters  
15 between parallel and serial ATA disk drive signals, (5) output connectors coupled to the  
16 converters, and (6) second switching logic capable of coupling a selected converter to a  
17 selected output connector. The second switching logic might also include hardware,  
18 firmware, software, or a combination or mixture thereof, and might be programmable.  
19 The second switching logic might also include control of a switched signal capable of  
20 supplying power to the disk drive.

## BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 shows a block diagram of a system including serial or parallel ATA disk drives, an adaptor capable of coupling to them and capable of coupling to a fiber channel backplane, and connectors capable of coupling multiple fiber channel backplanes to multiple fiber channel loops.

Figure 2 shows a state transition diagram used in a system for controlling programmable first or second switching logic as described above.

## INCORPORATED DISCLOSURE

Inventions described herein can be used in conjunction with technology described in the following documents:

- U.S. Patent Application Serial No. 09/860,027, filed May 16, 2001, in the name of David BULFER, Steven KLEIMAN, and Keith SON, titled "Reliable Coupling to an Array of Disk Storage Devices Using Single Initiator, Non-Hotplug Interfaces", attorney docket number 103.1058.01, client reference number P00-1058.01, and applications claiming priority therefrom.

1           These documents are hereby incorporated by reference as if fully set forth  
2 herein, and are sometimes referred to herein as the "incorporated disclosure."

3  
4           This application claims priority, to the fullest extent permitted by law,  
5 from the incorporated disclosure.

6  
7           DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

8  
9           In the description herein, a preferred embodiment of the invention is de-  
10 scribed, including preferred process steps and data structures. Those skilled in the art  
11 would realize, after perusal of this application, that embodiments of the invention  
12 might be implemented using a variety of other techniques not specifically described,  
13 without undue experimentation or further invention, and that such other techniques  
14 would be within the scope and spirit of the invention.

15  
16       *Lexicography*

17  
18           The general meaning of each of these following terms is intended to be il-  
19 lustrative and in no way limiting.

- 20  
21       • The phrases "mass storage," "mass storage system," and similar phrases and  
22 terms, generally describe systems in which data can be maintained, stored or



1       retrieved. As used herein, these phrases are intentionally broad. In the context  
2       of the invention, there is no particular requirement that “mass storage” must be  
3       relatively large or heavy.

- 4
- 5       • The phrases “disk drive,” “disk drive subsystem,” and similar phrases and  
6       terms, generally describe systems in which data can be maintained, stored or  
7       retrieved, such as for example systems using spinning disks and magnetic,  
8       optical, or a combination of magnetic and optical techniques. As used herein,  
9       these phrases are intentionally broad. Although one embodiment uses a set of  
10      individual disk drives, in the context of the invention, there is no particular  
11      requirement for using individual disk drives.

- 12
- 13     • The phrase “access point,” and similar phrases and terms, as applied to a disk  
14     drive, disk drive subsystem, or mass storage system, generally describe elements  
15     of a disk drivem, disk drive system, or mass storage system, from which data  
16     maintained thereon might be stored, retrieved, or manipulated. As used herein,  
17     this phrase is intentionally broad, and might include a computer system, file  
18     server, caching device, or other system used for accessing mass storage.

- 19
- 20     • The phrase “system for accessing,” and similar phrases and terms, as applied to a  
21     disk drive, disk drive subsystem, or mass storage system, describes any device  
22     for giving instructions to, making requests of, sending data to, or receiving data

or responses from, any mass storage system. As used herein, this phrase is intentionally broad, and might include a computer system, file server, caching device, or other system used for accessing mass storage. As used herein, this phrase is intentionally broad, and might include anything as described by the phrases "controller," "controlling device," "controlling system," and similar phrases.

- The phrases "multiple connectivity," "multiple coupling," "redundant capability," "redundant connectivity," "redundant coupling," "robust connectivity," "robust multiple connectivity," and similar phrases and terms, generally describe systems in which there is the possibility of more than one connection from an external point to stored data. As used herein, these phrases are intentionally broad, and might include any technique for providing failover, load balancing, clustering, high availability, hot swapping, or other techniques used for accessing mass storage using multiple connections.
- The phrase "independent interface," and similar phrases and terms, as applied to a disk drive, disk drive subsystem, or mass storage system, generally describe portions of systems by which data can be accessed without use of other interfaces. As used herein, these phrases are intentionally broad, and might include any technique for providing multiple connections between that mass storage system and one or more systems for accessing that mass storage system.

- 1
- 2     • The phrases “backplane,” “fiber channel backplane,” and similar phrases and
- 3       terms, generally describe portions of systems in which data could be sent,
- 4       received, or otherwise transferred among elements coupled to a common
- 5       medium, such as for example a backplane or fiber channel backplane. As used
- 6       herein, these phrases are intentionally broad, and include other types of
- 7       backplane interface besides standards known by the phrase “fiber channel” and
- 8       similar phrases.
- 9
- 10    • The phrase “form factor” generally describes an aspect of a system relating to
- 11      physical dimensions and shape, such as for example a shape of a disk drive
- 12      housing that could be fit into a disk drive housing holder. As used herein, this
- 13      phrase is intentionally broad.
- 14
- 15    • The phrases “logical switch,” “switching logic,” and similar phrases and terms,
- 16      generally describe techniques by which a choice of more than one signal could be
- 17      made, or by which more than one signal could be combined into one signal. As
- 18      used herein, these phrases are intentionally broad.
- 19
- 20    • The phrases “fiber channel loop,” and similar phrases and terms, generally
- 21      describe interfaces among devices using fiber channel communication
- 22      techniques. As used herein, these phrases are intentionally broad, and include

1 other types of backplane interface besides standards known by the phrase "fiber  
2 channel" and similar phrases.  
3

4 Other and further applications of the invention, including extensions of  
5 these terms and concepts, would be clear to those of ordinary skill in the art after pur-  
6 chasing this application. These other and further applications are part of the scope and  
7 spirit of the invention, and would be clear to those of ordinary skill in the art without  
8 further invention or undue experimentation.  
9

10 The scope and spirit of the invention is not limited to any of these defini-  
11 tions, or to specific examples mentioned therein, but is intended to include the most  
12 general concepts embodied by these and other terms.  
13

#### 14 *System Elements* 15

16 Figure 1 shows a block diagram of a system including serial or parallel  
17 ATA disk drives, an adaptor capable of coupling to them and capable of coupling to a  
18 fiber channel backplane, and connectors capable of coupling multiple fiber channel  
19 backplanes to multiple fiber channel loops.  
20

1           A system 100 includes a set of disk drive housings 110, a set of fiber chan-  
2 nel backplanes 120, a set of backplane couplers 130, a set of fiber channel loops 140, and  
3 a set of controlling devices 150.

#### 4 5           *Disk Drive Housings*

6  
7           Each disk drive housing 110 includes a disk drive 111, having a disk drive  
8 interface 112 by which external devices communicate with the disk drive 111. In one  
9 embodiment, as described herein, the disk drive 111 might include an ATA disk drive,  
10 such as for example a P-ATA (parallel ATA) disk drive or an S-ATA (serial ATA) disk  
11 drive, that is, a disk drive 111 with a disk drive interface 112 responding according to  
12 known specifications for P-ATA or S-ATA disk drives. In other embodiments, as also  
13 described herein, the disk drive 111 might include an FC (fiber channel) disk drive, that  
14 is, a disk drive 111 with disk drive interface 112 responding according to known specifi-  
15 cations for FC disk drives.

16  
17           Each disk drive housing 110 has a physical form factor and a set of electri-  
18 cal connections compatible with a known specification for FC disk drives. This has the  
19 effects that (a) each disk drive housing 110 can be physically and electrically coupled to  
20 one of the fiber channel backplanes 120, and (b) each disk drive housing 110, being  
21 physically compatible with the fiber channel backplanes 120, does not impose any ad-  
22 ditional physical stress on connectors between that disk drive housing 110 and that fiber

1 channel backplane 120, than might be imposed if there were more than one physical  
2 form factor used for the disk drive housings 110. However, as further described below,  
3 while each disk drive housing 110 has a physical form factor compatible with a known  
4 specification for FC disk drives, in the context of the invention, there is no particular re-  
5 quirement that any of the disk drive housings 110 are electrically compatible with  
6 known specifications for FC disk drives, and in fact, in one embodiment, at least some  
7 of the disk drive housings 110 use a set of electrical signals possibly incompatible with  
8 known specifications for FC disk drives.

9  
10 After reading this application, those skilled in the art would recognize that  
11 disparate disk drive housings 110, having distinct types, might be coupled to the same  
12 fiber channel backplane 120. Those disk drive housings 110 might have different types  
13 of disk drives 111 or different types of disk drive interfaces 112.

- 14
- 15 • For a first example, some of those disk drive housings 110 might have P-ATA  
16 disk drives 111, while others of those disk drive housings 110 might have S-ATA  
17 disk drives 111, as further described below.

- 18
- 19 • For a second example, some of those disk drive housings 110 might have ATA  
20 disk drives 111, while others of those disk drive housings 110 might have FC  
21 disk drives 111.

- For a third example, some of those disk drive housings 110 might have more involved disk drive subsystems or mass storage, such as for example RAID disk drive subsystems or other types of highly-available, redundant, or otherwise intelligently controlled disk drive subsystems or mass storage.

Those disk drive housings 110 having ATA disk drives 111 include an adaptor 160 coupled to the disk drive interface 112, with the effects of (a) being capable of receiving signals from the disk drive interface 112 and (b) being capable of sending signals to the disk drive interface 112.

#### *Switching Adaptor*

The adaptor 160 includes a first communication path 161 and a second communication path 162, each of which is coupled to the disk drive interface 112 using a known disk drive connector 163 (such as for example a P-ATA connector or an S-ATA connector). The first communication path 161 and the second communication path 162 are each coupled to a backplane connector 164 having a physical form factor compatible with the fiber channel backplane 120.

In one embodiment, of the first communication path 161 and the second communication path 162, either one or both might include a serial-to-parallel converter or a parallel-to-serial converter (or neither), with the effect that P-ATA disk drives 111

1 and S-ATA disk drives 111 are capable of inter-communicating using the fiber channel  
2 bus 120 and one or more adaptors 160. Thus, the disk drive connector 163 might be ei-  
3 ther a P-ATA connector or an S-ATA connector, with the effect that P-ATA disk drives  
4 111 and S-ATA disk drives 111 are each capable of being coupled to the adaptor 160.

5  
6 However, in the context of the invention, there is no requirement the  
7 backplane connector 164 uses standard electrical signals for a fiber channel backplane  
8 120 (such as for example the specification for fiber channel backplane 120 electrical sig-  
9 nals found in the SFF-8045 document, available on the web via [www.sff.org](http://www.sff.org), hereby in-  
10 corporated by reference as if fully set forth herein). Rather, in one set of embodiments,  
11 the backplane connector 164 uses a set of electrical signals not standard to a fiber chan-  
12 nel backplane 120, but instead capable of coupling both the first communication path  
13 161 and the second communication path 162 to separate signal communication lines on  
14 the fiber channel backplane 120. In such embodiments, the fiber channel backplane 120  
15 is disposed to communicate a first data signal "A" (coupled to the first communication  
16 path 161), and a set of control signals associated therewith, and a like set of signals for a  
17 second data signal "B" (coupled to the second communication path 162).

18  
19 For example, the control signals associated with each of the first data sig-  
20 nal "A" and the second data signal "B" might include:



- 1       • a set of transmit and receive signals (such as for example those sometimes  
2       known as Tx +/- and Rx +/-);  
3
- 4       • a set of power and ground signals (such as for example a +12 volt power signal, a  
5       +5 volt power signal, a grounded power signal, and a ground-to-frame power  
6       signal); and  
7
- 8       • a set of other known control signals.  
9

10           The adaptor 160 includes a logical switch 165 coupled to the first commu-  
11   nication path 161 and the second communication path 162, and capable of disabling or  
12   enabling the first communication path 161 or the second communication path 162, with  
13   the effect of selecting which one is coupled between the disk drive connector 163 and  
14   the backplane connector 164. In one embodiment, the switch 165 might include hard-  
15   ware, firmware, or software, and might be programmable, with the effect that the  
16   switch 165 is capable of making a switching decision in response to information on the  
17   first communication path 161 or the second communication path 162.  
18

19           In such embodiments, the switch 165 is coupled to an external switching  
20   signal 166. The external switching signal 166 might include information for an explicit  
21   switching decision, or might include instructions capable of being interpreted by the  
22   switch 165 (or a computing device located therein, or coupled thereto) for making a

1 switching decision. In one embodiment, the switch 165 is pre-selected with known in-  
2 structions, with the effect of following a state diagram such as shown in figure 2 below.

### 4 *Fiber Channel Backplane*

5  
6 Each fiber channel backplane 120 includes a backplane compatible with a  
7 known fiber channel standard, such as for example the specification for a fiber channel  
8 backplane, its physical form factor and its electrical signals, found in the SFF-8045  
9 document described and incorporated above, and sometimes referred to herein as the  
10 "fiber channel standard."

11  
12 The fiber channel backplane 120 includes a backplane communication link  
13 121 compatible with the same fiber channel standard. As further described in the fiber  
14 channel standard, the backplane communication link 121 includes forty control signals  
15 suitable for communication among FC disk drives 111. However, as further described  
16 in this application, in one embodiment the backplane 121 includes two sets of forty  
17 control signals (multiplexed 2:1 by the logical switch 165 to forty control signals), suit-  
18 able for communication among ATA disk drives 111, as mediated by the adaptor 160  
19 and its logical switch 165.

20  
21 The fiber channel backplane 120 includes a backplane controller 122 com-  
22 patible with same fiber channel standard. As further described in the fiber channel

1 standard, the backplane controller 122 is capable of sending and receiving communica-  
2 tions between disk drives 111 and the fiber channel backplane 120. As further described  
3 in that document, the backplane controller 122 is capable of sending and receiving  
4 communications between the fiber channel backplane 120 and one or more fiber chan-  
5 nel loops 140. However, as further described in this application, in one embodiment the  
6 backplane controller 122 includes a backplane coupler 130, capable of sending and re-  
7 ceiving signals between the ATA disk drives 111 and the one or more fiber channel  
8 loops 140.

#### 9 10 *Backplane Coupler*

11  
12 Each fiber channel backplane 120 having ATA disk drives 111 coupled to  
13 it also has a backplane coupler 130 coupled to it, and capable of sending and receiving  
14 signals between the fiber channel backplane 120 and the one or more fiber channel  
15 loops 140.

16  
17 Each backplane coupler 130 includes fiber channel translation device 137  
18 coupled using one of its backplane communication links 121 to the first communication  
19 path 161 and the second communication path 162 at one of the adaptors 160. The fiber  
20 channel translation device 137 receives the ATA signals from the fiber channel bus 130,  
21 and translates them to fiber channel standard signals. The fiber channel translation de-  
22 vice 137 is coupled to a processor 136 and a memory 135, each of which are coupled to

each other and to a PCI bus 134. The PCI bus 134 is coupled to a first communication port 131 and a second communication port 132.

After reading this application, those skilled in the art will recognize that the PCI bus 134, the memory 135, and the processor 136, each operate as such known devices, with the processor 136 being under control of instructions it interprets to perform functions involved in translating between ATA-compatible signals and FC-compatible signals. The fiber channel translation device 137 sends and receives signals between the memory 135 and the processor 136 (on the one hand), and the backplane communication link 121 (on the other hand).

#### *Fiber Channel Loops*

Each fiber channel loop 140 couples a set of fiber channel backplanes 120 in series, with the effect that each disk drive 111 coupled to at least one of the fiber channel backplanes 120 is capable of communication with each controlling device 150 coupled to that fiber channel loop 140.

In one embodiment, more than one such fiber channel loop 140 is available for communication, with the effect that each controlling device 150 coupled to those fiber channel loops 140 is capable of communication with each disk drive 111

1 coupled to at least one of the fiber channel backplanes 120 coupled to those fiber chan-  
2 nel loops 140.

### 4 *Controlling Devices*

5  
6 As described above, the controlling devices 150 might include any system  
7 for accessing one or more of the disk drives 111 or disk drive interfaces 112, and might  
8 include for example any computer system, file server, caching device, or other system  
9 used for accessing mass storage. For example, any device having the capability of being  
10 a controller, controlling device, or controlling system, for disk drives or other mass  
11 storage, might be included within the capability of one or more of the controlling de-  
12 vices 150. Similarly, any device having the function of providing access to disk drives  
13 or other mass storage, such as for example, any file system, filer, FAS (fabric attached  
14 storage), JBOD ("just a bunch of disks"), LUN (local unit network), NAS (network at-  
15 tached storage), SAN (storage area network), SAN/NAS appliance, virtual disk, data-  
16 base, data container, data repository, or any other device or system having the capabil-  
17 ity of receiving and responding to requests relating to storage, retrieval, or maintenance  
18 of data, might be included within the capability of one or more of the controlling de-  
19 vices 150.

20  
21 After reading this application, those skilled in the art would recognize that  
22 each such controlling device 150 has more than one path by which it might access each

1 of the disk drives 111 or disk drive interfaces 112, with the effect that each of the disk  
2 drives 111 or disk drive interfaces 112 is described herein as having more than one ac-  
3 cess point or independent interface, and with the effect that each of the disk drives 111  
4 or disk drive interfaces 112 is described herein as having multiple connectivity (to the  
5 controlling devices 150), multiple coupling thereto, redundant capability (for being ac-  
6 cessed thereby), redundant connectivity thereto, redundant coupling thereto, robust  
7 connectivity thereto, or robust multiple connectivity thereto.

8  
9           After reading this application, those skilled in the art would recognize that  
10 each such controlling device 150 with multiple connectivity to disk drives 111 (or disk  
11 drive interfaces 112) is capable of performing known techniques in any type of redun-  
12 dant storage access structure for redundant and robust use of storage on those disk  
13 drives 111. These known techniques include failover, load balancing, clustering, high  
14 availability, hot swapping, and other advantages allowed by multiple controlling de-  
15 vices 150 interfacing with multiple disk drives 111 (or disk drive interfaces 112).

16  
17 *State Transition Diagram*

18  
19           Figure 2 shows a state transition diagram used in a system for controlling  
20 programmable first or second switching logic as described above.

1           A state transition diagram 200 includes a set of states 201 (Idle), 202 (A  
2 Owns), and 203 (B Owns), a set of transitions for those states, subject to signals relating  
3 to those transitions.

4  
5           A transition 211, responsive to a “reset” signal, causes the system to enter  
6 the state 201 (Idle), that is, neither the first communication path 161 (“A”) nor the sec-  
7 ond communication path 162 (“B”) controls the connection between the disk drive inter-  
8 face 112 and the fiber channel backplane 120.

9  
10          A transition 212, responsive to an “A” signal (that is, A desires to own the  
11 connection) and to a “B” signal (that is, B desires to own the connection), causes the  
12 system from the state 201 (Idle) to re-enter the state 201 (Idle) if the value (!A & !B) is  
13 logically true, where !X indicates negation of a logical value X.

14  
15          A transition 213 causes the system from the state 201 (Idle) to enter the  
16 state 202 (A Owns) if the value (A) is logically true.

17  
18          A transition 214 causes the system from the state 201 (Idle) to enter the  
19 state 203 (B Owns) if the value (!A & B) is logically true.

20  
21          A transition 215 causes the system from the state 202 (A Owns) to enter  
22 the state 201 (Idle) if the value (!A) is logically true.

1  
2 A transition 216, responsive to the A signal and the B signal, and to a  
3 "B\_Delay" signal indicative of a transmission delay on the second communication path  
4 162 ("B"), causes the system from the state 202 (A Owns) to re-enter the state 202 (A  
5 Owns) if the value  $((A \ \& \ !B) \text{ or } (A \ \& \ B \ \& \ B\_Delay))$  is logically true.

6  
7 A transition 217 causes the system from the state 202 (A Owns) to enter  
8 the state 203 (B Owns) if the value  $(A \ \& \ B \ \& \ !B\_Delay)$  is logically true.

9  
10 A transition 218 causes the system from the state 203 (B Owns) to enter the  
11 state 201 (Idle) if the value  $(!B)$  is logically true.

12  
13 A transition 219, responsive to the A signal and the B signal, and to an  
14 "A\_Delay" signal indicative of a transmission delay on the first communication path  
15 161 ("A"), causes the system from the state 203 (B Owns) to re-enter the state 203 (B  
16 Owns) if the value  $((!A \ \& \ B) \text{ or } (A \ \& \ B \ \& \ A\_Delay))$  is logically true.

17  
18 A transition 220 causes the system from the state 203 (B Owns) to enter the  
19 state 202 (A Owns) if the value  $(A \ \& \ B \ \& \ !A\_Delay)$  is logically true.

20  
21 After reading this application, those skilled in the art would understand  
22 that other state transition diagrams 200 could also be suitable for controlling the system,



1 that such other state transition diagrams 200 are within the scope and spirit of the in-  
2 vention, and that such other state transition diagrams 200 could be implemented with-  
3 out either undue experimentation or further invention.

4  
5 *Alternative Embodiments*

6  
7 Although preferred embodiments are disclosed herein, many variations  
8 are possible which remain within the concept, scope, and spirit of the invention. These  
9 variations would become clear to those skilled in the art after perusal of this applica-  
10 tion.

- 11
- 12 • The invention applies to any mass storage device or system, and is not limited to  
13 disk drives or their equivalents.
  - 14
  - 15 • The invention applies to any form of mass storage controller, and is not limited  
16 to S-ATA, P-ATA, or fiber channel disk drives.
  - 17
  - 18 • The invention applies to any form of adaptor between mass storage device or  
19 system form factors, and is not limited to S-ATA, P-ATA, or fiber channel disk  
20 drives.
  - 21

- 1       • The invention applies to systems including any form of switch, and is not limited  
2       to the programmable switches described herein.  
3
- 4       • The invention applies to systems having multiple connectivity with relatively  
5       complete redundancy, to systems having multiple connectivity but without such  
6       redundancy, and to systems not having multiple connectivity.  
7
- 8       • The invention applies to systems having any form of inter-disk connectivity, and  
9       is not limited to fiber channel backplanes or to fiber channel loops.  
10
- 11      • The invention applies to systems in which the disk drives connected to a single  
12      backplane are heterogeneous, and to systems in which the disk drives connected  
13      to a single backplane are not heterogeneous.  
14
- 15      • The invention applies to systems in which signals are converted between any  
16      types of disk drives, controllers, or backplanes, and is not limited to conversion  
17      between ATA and fiber channel signals.  
18
- 19      • The invention applies to systems including techniques for failover, load balanc-  
20      ing, clustering, high availability, and hot swapping, and to systems not including  
21      such techniques.  
22

- 1           After reading this application, those skilled in the art will recognize that
- 2   these alternative embodiments are illustrative and in no way limiting.